ATTACHMENT D

Network address translation

From Wikipedia, the free encyclopedia

In computer networking, **network address translation** (NAT, also known as *network masquerading*, *native address translation* or *IP masquerading*) is a technique of transceiving network traffic through a router that involves re-writing the source and/or destination IP addresses and usually also the TCP/UDP port numbers of IP packets as they pass through. Checksums (both IP and TCP/UDP) must also be rewritten to take account of the changes. Most systems using NAT do so in order to enable multiple hosts on a private network to access the Internet using a single public IP address (see gateway). Nonetheless, NAT can introduce complications in communication between hosts and may have a performance impact.

Contents

Overview

NAT first became popular as a way to deal with the IPv4 address shortage and to avoid all the difficulty of reserving IP addresses. NAT has proven particularly popular in countries other than the United States, which (for historical reasons) have fewer address-blocks allocated per capita. It has become a standard feature in routers for home and small-office Internet connections, where the price of extra IP addresses would often outweigh the benefits. NAT also adds to security as it disguises the internal network's structure: all traffic appears to outside parties as if it originates from the gateway machine.

In a typical configuration, a local network uses one of the designated "private" IP address subnets (the RFC 1918 Private Network Addresses are 192.168.x.x, 172.16.x.x through 172.31.x.x, and 10.x.x.x - using CIDR notation, 192.168/16, 172.16/12, and 10/8), and a router on that network has a private address (such as 192.168.0.1) in that address space. The router is also connected to the Internet with a single "public" address (known as "overloaded" NAT) or multiple "public" addresses assigned by an ISP. As traffic passes from the local network to the Internet, the source address in each packet is translated on the fly from the private addresses to the public address(es). The router tracks basic data about each active connection (particularly the destination address and port). When a reply returns to the router, it uses the connection tracking data it stored during the outbound phase to determine where on the internal network to forward the reply; the TCP or UDP client port numbers are used to demultiplex the packets in the case of overloaded NAT, or IP address and port number when multiple public addresses are available, on packet return. To a system on the Internet, the router itself appears to be the source/destination for this traffic.

It has been argued that the wide adoption of IPv6 would make NAT unnecessary, as NAT is a method of handling the shortage of IPv4 address space.

Drawbacks

Hosts behind NAT-enabled routers do not have true end-to-end connectivity and cannot participate in some Internet protocols. Services that require the initiation of TCP connections from the outside network, or stateless protocols such as those using UDP, can be disrupted. Unless the NAT router makes a specific effort to support such protocols, incoming packets cannot reach their destination. Some protocols can accommodate one instance of NAT between participating hosts ("passive mode" FTP, for example), sometimes with the assistance of an Application Layer Gateway (see below), but fail when both systems are separated from the Internet by NAT. Use of NAT also complicates tunneling protocols such as IPsec because NAT modifies values in the headers which interfere with the integrity checks done by IPsec and other tunneling protocols.

End-to-end connectivity has been a core principle of the Internet, supported for example by the Internet Architecture Board. Current Internet architectural documents observe that NAT is a violation of the End-to-End Principle, but that NAT does have a valid role in careful design.[1] There is considerably more concern with the use of IPv6 NAT, and many IPv6 architects believe IPv6 was intended to remove the need for NAT.[2]

Some Internet service providers (ISPs) only provide their customers with "local" IP addresses.Thus, these customers must access services external to the ISP's network through NAT. As a result, it may be argued that such companies do not properly provide "Internet" service.

Benefits

In addition to the convenience and low cost of NAT, the lack of full bidirectional connectivity can be regarded in some situations as a feature rather than a limitation. To the extent that NAT depends on a machine on the local network to initiate any connection to hosts on the other side of the router, it prevents malicious activity initiated by outside hosts from reaching those local hosts. This can enhance the reliability of local systems by stopping worms and enhance privacy by discouraging scans. Many NAT-enabled firewalls use this as the core of the protection they provide.

The greatest benefit of NAT is that it is a practical solution to the impending exhaustion of IPv4 address space. Networks that previously required a Class B IP range or a block of Class C network addresses can now be connected to the Internet with as little as a single IP address (many home networks are set up this way). The more common arrangement is having machines that require true bidirectional and unfettered connectivity supplied with a 'real' IP address, while having machines that do not provide services to outside users tucked away behind NAT with only a few IP addresses used to enable Internet access.

Basic NAT and PAT

Two kinds of network address translation exist:

- **PAT (Port Address Translation)** - The type popularly, but incorrectly, called simply "NAT" (also sometimes named "Network Address Port Translation, NAPT") refers to network address translation involving the mapping of port numbers, allowing multiple machines to share a single IP address.
- **Basic NAT** - The other, technically simpler, forms—"one-to-one NAT", "basic NAT", "static NAT" and "pooled NAT"—involve only address translation, not port mapping. This requires an external IP address for each simultaneous connection. Broadband routers often use this feature, sometimes labelled "DMZ host", to allow a designated computer to accept all external connections even when the router itself uses the only available external IP address.

NAT with port-translation (i.e. PAT) comes in two sub-types: source address translation (**source NAT**), which re-writes the IP address of the computer which initiated the connection; and its counterpart, destination address translation (**destination NAT**). In practice, both are usually used together in coordination for two-way communication.

NOTE: 'PAT' as it is referred to here is referred to by Cisco as NAT 'overloading' as described in this Howstuffworks article provided to Howstuffworks by Cisco: http://computer.howstuffworks.com/nat3.htm

Relationship between NAT and PAT

PAT is closely related to NAT.

In NAT, generally only the IP addresses are modified: There is a 1:1 correspondence between publicly exposed IP addresses and privately held IP addresses. In PAT, both the sender's private IP and port number are modified; the PAT device chooses the port numbers that will be seen by hosts on the public network.

In NAT, incoming packets are routed to their destination IP address on the private network by reference to the incoming source IP address given by the host on the public network. In PAT, there is generally only one publicly exposed IP address and incoming packets from the public network are routed to their destinations on the private network by reference to a table held within the PAT device that keeps track of public and private port pairs. This is often called connection tracking.

Some devices that claim to offer NAT, such as broadband routers, actually offer PAT. For this reason, there is considerable confusion between the terms. The common use of NAT to include PAT devices suggests that PAT should be considered a type of NAT rather than a distinct technology.

NAT and TCP/UDP

"Pure NAT", operating on IP alone, may or may not correctly pass protocols that are totally concerned with IP information, such as ICMP, depending on whether the payload is interpreted by a host on the "inside" or "outside" of translation. As soon as the protocol stack is climbed, even with such basic protocols such TCP and UDP, the protocols will break unless NAT takes action beyond the network layer.

IP has a checksum in each packet header, which provides error detection only for the header. IP datagrams may become fragmented and it is necessary for a NAT to reassemble these fragments to allow correct recalculation of higher level checksums and correct tracking of which packets belong to which connection.

The major transport layer protocols, TCP and UDP, have a checksum that covers all the data they carry, as well as the TCP/UDP header, plus a "pseudo-header" that contains the source and destination IP addresses of the packet carrying the TCP/UDP header. For an originating NAT to successfully pass TCP or UDP, it must recompute the TCP/UDP header checksum based on the translated IP addresses, not the original ones, and put that checksum into the TCP/UDP header of the first packet of the fragmented set of packets. The receiving NAT must recompute the IP checksum on every packet it passes to the destination host, and also recognize and recompute the TCP/UDP header using the retranslated addresses and pseudo-header. This is not a completely solved problem. One solution is for the receiving NAT to reassemble the entire segment and then recompute a checksum calculated across all packets.

It may be wise for the originating host to do MTU Path Discovery (RFC1191) to determine what MTU will go to the end without fragmentation, and then set the "don't fragment" bit in the appropriate packets. There is no totally general solution to this problem, which is why one of the goals of IPv6 is to avoid NAT.

Applications affected by NAT

Some higher-layer protocols (such as FTP and SIP) send network layer address information inside application payloads. FTP in active mode, for example, uses separate connections for control traffic (commands) and for data traffic (file contents). When requesting a file transfer, the host making the request identifies the corresponding data connection by its network layer and transport layer addresses. If the host making the request lies behind a simple NAT firewall, the translation of the IP address and/or TCP port number makes the information received by the server invalid.

An Application Layer Gateway (ALG) can fix this problem. An ALG software module running on a NAT firewall device updates any payload data made invalid by address translation. ALGs obviously need to understand the higher-layer protocol that they need to fix, and so each protocol with this problem requires a separate ALG.

Another possible solution to this problem is to use NAT traversal techniques using protocols such as STUN or ICE or proprietary approaches in a session border controller. NAT traversal is possible in both TCP- and UDP-based applications, but the UDP-based technique is simpler, more widely understood, and more compatible with legacy NATs. In either case, the high level protocol must be designed with NAT traversal in mind, and it does not work reliably across symmetric NATs or other poorly-behaved legacy NATs.

Other possibilities are UPnP (Universal Plug and Play) or Bonjour (NAT-PMP), but these require the cooperation of the NAT device.

Most traditional client-server protocols (FTP being the main exception), however, do not send layer 3 contact information and therefore do not require any special treatment by NATs. In fact, avoiding NAT complications is practically a requirement when designing new higher-layer protocols today.

NATs can also cause problems where IPsec encryption is applied and in cases where multiple devices such as SIP phones are located behind a NAT. Phones which encrypt their signalling with IPsec encapsulate the port information within the IPsec packet meaning that NA(P)T devices cannot access and translate the port. In these cases the NA(P)T devices revert to simple NAT operation. This means that all traffic returning to the NAT will be mapped onto one client causing the service to fail. There are a couple of solutions to this problem, one is to use TLS which operates at level 4 in the OSI Reference Model and therefore does not mask the port number, or to Encapsulate the IPsec within UDP - the latter being the solution chosen by TISPAN to achieve secure NAT traversal.

It is also a problem with Xbox Live gameplay when the need for players to communicate becomes harder and slows down gameplay.

Different types of NAT

Applications that deal with NAT sometimes need to characterize NAT by type. The STUN protocol proposed to characterize Network address translation as **Full cone NAT**, **restricted cone NAT**, **port restricted cone NAT** or **symmetric NAT**.[3][4]

Full cone NAT, also known as one-to-one NAT ▪ Once an internal address (iAddr:port1) is mapped to an external address (eAddr:port2), any packets from iAddr:port1 will be sent through eAddr:port2. Any external host can send packets to iAddr:port1 by sending packets to eAddr:port2.	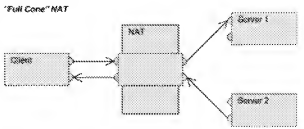
Restricted cone NAT ▪ Once an internal address (iAddr:port1) is mapped to an external address (eAddr:port2), any packets from iAddr:port1 will be sent through eAddr:port2. An external host (hostAddr:any) can send packets to iAddr:port1 by sending packets to eAddr:port2 only if iAddr:port1 had previously sent a packet to hostAddr:any. "any" means the port number doesn't matter.	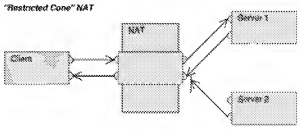
Port restricted cone NAT Like a **restricted cone NAT**, but the restriction includes port numbers. ▪ Once an internal address (iAddr:port1) is mapped to an external address (eAddr:port2), any packets from iAddr:port1 will be sent through eAddr:port2. An external host (hostAddr:port3) can send packets to iAddr:port1 by sending packets to eAddr:port2 only if iAddr:port1 had previously sent a packet to hostAddr:port3.	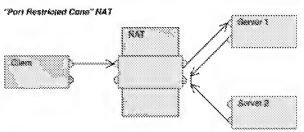
Symmetric NAT ▪ Each request from the same internal IP address and port to a specific destination IP address and port is mapped to a unique external source IP address and port. If the same internal host sends a packet even with the same source address and port but to a different destination, a different mapping is used. ▪ Only an external host that receives a packet from an internal host can send a packet back.	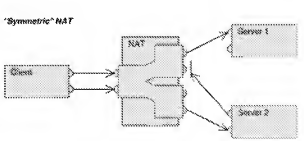

This terminology has been the source of much confusion, as it has proven inadequate at describing real-life NAT behavior.[5] Many NAT implementations combine the specified types, and it is therefore better to refer to specific individual NAT behaviors instead of using the Cone/Symmetric terminology. Especially, most NAT translators combine *symmetric NAT* for outgoing connections with a *static port mapping* capability. The latter means that all incoming packets to the specific external address and port can be redirected to a specific internal address and port. Some products can redirect packets to several internal hosts - e.g. to divide the load between a few servers (however, this introduces problems with more sophisticated communications having many interconnected packets and thus is rarely used.

Many NAT implementations follow a **port preservation** design. For most communications, they will use the same values as internal and external port numbers. However, if two internal hosts attempt to communicate with the same external host using the same port number, the external port number used by the second host will be chosen at random. Such NAT will be sometimes perceived as **restricted cone NAT** and other times as **symmetric NAT**.

Other examples of use

- Load Balancing: Destination NAT can redirect connections pointed at some server to randomly selected servers.
- Failover: Destination NAT can be used to set up a service requiring high availability. If a system involves a critical server accessed through a router, and if the router detects that server has gone down, it could use destination NAT to transparently re-route a connection to arrive on a backup server.
- Transparent proxying: NAT can redirect HTTP connections targeted at the Internet to a special HTTP proxy which can cache content and filter requests. Some internet service providers use this technique to reduce bandwidth usage without requiring their clients to configure their web browser for proxy support.
- Overlapping Networks: Advanced NAT configurations can connect two networks that have overlapping addresses, such as Private addresses, and this is a very useful feature for companies that have just started to merge their networks. This requires that both Source & Destination IP addresses be replaced at the same time, fooling each other's networks.

See also

- Port address translation (PAT, can be used in conjunction with NAT)
- Firewall (networking)
- Proxy server
- Middlebox
- Routing
- Gateway (telecommunications)
- Subnet
- IPv6
- UDP hole punching
- AYIYA (IPv6 over IPv4 UDP thus working IPv6 tunneling over most NATs)
- IPv4 address exhaustion
- Private IP address
- Internet Connection Sharing
- Port forwarding
- STUN: Simple Traversal of UDP over NATs
- Teredo tunneling: NAT traversal using IPv6
- Internet Gateway Device (IGD) Protocol: UPnP NAT-traversal method
- NAT-PT

Popular NAT software

- IPFilter
- PF (firewall): The OpenBSD Packet Filter.
- Iptables masquerading
- Berkeley Software Distribution
- Internet Connection Sharing (ICS)
- WinGate
- Cisco IOS

References

1. ^ Some Internet Architectural Guidelines and Philosophy,RFC 3439, R. Bush & D. Meyer,December 2002
2. ^ Local Network Protection for IPv6,RFC 4864, G. Van de Velde *et al.*,May 2007
3. ^ STUN
4. ^ NAT Types (PDF).
5. ^ Francois Audet, Cullen Jennings (January 2007). "*RFC 4787 Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*" (text). IETF. Retrieved on 2007-08-29.

External links

- TeleDict: Introduction to NAT
- Characterization of different TCP NATs – Paper discussing the different types of NAT
- P2P & NAT connections (In Spanish PDF)
- Test the NAT type of your routerA small tool by the eMule developers to detect the nat type of your router.
- Anatomy: A Look Inside Network Address Translators – Volume 7, Issue 3, September 2004
- HowStuffWorks: *How Network Address Translation Works* by Jeff Tyson
- NAT traversal techniques in multimedia Networks – White Paper from Newport Networks

- Peer-to-Peer Communication Across Network Address Translators (PDF) – NAT traversal techniques for UDP and TCP
- RFC 4008 – Standards Track – Definitions of Managed Objects for Network Address Translators (NAT)
- RFC 3022 – Traditional IP Network Address Translator (Traditional NAT)
- RFC 1631 – Obsolete – The IP Network Address Translator (NAT)
- nat-traverse – Tool to establish tunnels through NAT gateways without need for reconfiguration of the involved routers
- Speak Freely End of Life Announcement – John Walker's discussion of why he stopped developing a famous program for free internet communication, part of which is directly related to NAT.
- NATs, IPsec and VoIP – White paper looks at the issues of IPsec and NAT devices and how UDP encapsulation of IPsec solves the problems.
- Cisco IOS NAT Application Layer Gateway
- natd
- A blog entry explanation of NAT
- Alternative Taxonomy
  - Static NAT
  - Dynamic NAT
  - Masquerade NAT